

Docket No. 113605

OLIFF & BERRIDGE, PLC
P.O. Box 19928
Alexandria, Virginia 22320
Telephone: (703) 836-6400
Facsimile: (703) 836-2787
Assistant Commissioner for Patents
Washington, D.C. 20231

NONPROVISIONAL PATENT APPLICATION

Attorney Docket No.: Docket No. 113605

Date: October 19, 1999

BOX PATENT APPLICATION

NONPROVISIONAL APPLICATION TRANSMITTAL
RULE §1.53(b)

Sir:

Transmitted herewith for filing under 37 C.F.R. §1.53(b) is the nonprovisional patent application

For (Title): VIRTUAL PRIVATE NETWORK

Casename: Docket No. 113605

By (Inventors): Nicholas G. DUFFIELD; Albert G. GREENBERG; Pawan GOYAL; Partho P. MISHRA; Kadangode K. RAMAKRISHNAN; and Jacobus Erasmus van der MERWE

Formal drawings (Figs. 1-9; nine sheets) are attached.
 A Declaration and Power of Attorney is filed herewith.
 An assignment of the invention to _____ is filed herewith.
 An Information Disclosure Statement is filed herewith.
 A Preliminary Amendment is filed herewith.
 Please amend the specification by inserting before the first line the sentence --This nonprovisional application claims the benefit of U.S. Provisional Application No. _____, filed _____.--
 Priority of foreign application (s) No. _____ filed _____ in _____ is claimed under 35 U.S.C. §119.
 A certified copy of the above corresponding foreign application(s) is filed herewith.
 The filing fee is calculated below:

CLAIMS IN THE APPLICATION AFTER ENTRY
OF ANY PRELIMINARY AMENDMENT NOTED ABOVE

FOR:	NO. FILED	NO. EXTRA	RATE	FEE
BASIC FEE				\$ 760
TOTAL CLAIMS	39 - 20	= *	x 18	\$ 342
INDEP CLAIMS	2 - 3	= *	x 78	\$
<input type="checkbox"/> MULTIPLE DEPENDENT CLAIMS PRESENTED		+260		\$
* If the difference is less than zero, enter "0".			TOTAL	\$1102

Check No. 103743 in the amount of \$1102 to cover the filing fee is attached. Except as otherwise noted herein, the Commissioner is hereby authorized to charge any other fees that may be required to complete this filing, or to credit any overpayment, to Deposit Account No. 15-0461. Two duplicate copies of this sheet are attached.
 Please address all correspondence to S. H. Dworetzky, AT&T Corp., P.O. Box 4110, Middletown, New Jersey 07748. Telephone calls should be made to Mr. Dworetzky at Area Code (973) 360-8120.

Respectfully submitted,

James A. Oliff, Reg. No. 27,075
Paul Tsou, Reg. No. 37,956
Attorney for Applicants

JAO:PT/tlp
Date: October 19, 1999

OLIFF & BERRIDGE, PLC
P.O. Box 19928
Alexandria, Virginia 22320
Telephone: (703) 836-6400

DEPOSIT ACCOUNT USE
AUTHORIZATION
Please grant any extension
necessary for entry;
Charge any fee due to our
Deposit Account No. 15-0461

VIRTUAL PRIVATE NETWORK

BACKGROUND OF THE INVENTION

This non-provisional application claims the benefit of U.S. Provisional Application No. 60/113,497, entitled "PERFORMANCE ORIENTED SERVICE INTERFACE FOR VIRTUAL PRIVATE NETWORKS" filed on December 22, 1998 (Attorney Docket No. DUFFIELD 6) and U.S. Provisional Application No. 60/104,756, entitled "VIRTUAL PRIVATE NETWORK" filed on October 19, 1998 (Attorney Docket No. 113470). The Applicants of the provisional applications are Nicholas G. Duffield, Pawan Goyal, Albert Gorder Greenberg, Partho Pratim Mishra, Kadangode K. Ramakrishnan, and Jacobus E. van der Merwe. The above provisional applications and "A FLEXIBLE MODEL FOR RESOURCE MANAGEMENT IN VIRTUAL PRIVATE NETWORKS" by Duffield, Goyal, Greenberg, Mishra, Ramakrishnan and van der Merwe are hereby incorporated by reference including all references cited therein.

1. Field of Invention

The present invention relates to methods and apparatus for virtual private networks.

2. Description of Related Art

A private line network provides communications with guaranteed performance because entities that are connected to the private line network are limited to only those parties that communicate with one another. In addition, the private line network provides a level of inherent security because parties other than the communicating parties cannot gain access to the network.

The above advantages are counter balanced by inflexibility (e.g., adding new endpoints requires new network upgrades), and network resources are determined by peak traffic conditions. Moreover, interfacing to private line networks is complex because detailed knowledge about the network is required, including particularities regarding endpoints and connectivity to the endpoints. Thus, there is a need for new technology to provide private line network service without the above mentioned drawbacks.

SUMMARY OF THE INVENTION

The invention provides apparatus and methods for a Virtual Private Network (VPN) in a network such as an Internet Protocol (IP) Network that offers a simple user interface for efficient utilization of network resources. The VPN is defined for a specified set of endpoints, each of which is associated with a single "hose." A hose provides access to the VPN through an access point which may be a node of the network such as a server, a router or an Internet Service Provider (ISP), for example. The hose is a single interface for a user or "customer" to the VPN for communication to all other endpoints of the VPN.

Customers of the VPN are relieved of the burden to obtain detailed knowledge of the VPN. To interface with the VPN, each customer is only required to specify a hose profile of the hose associated with the customer's endpoint. The hose profile may be an aggregate communication specification such as bandwidth or other quality of service requirements or more detailed description of the desired communication traffic characteristics. The VPN guarantees communication performance up to the hose profile for each of the endpoints.

The VPN achieves network resource allocation efficiency by exploiting resource sharing possibilities via multiplexing routing paths between endpoints and dynamic resource allocation techniques that permit real time resource allocation resizing. When a VPN is established with a VPN service provider, the routing paths between the endpoints of the VPN is optimized for multiplexing opportunities so that resource allocations between nodes along routing paths within the IP network is reduced to a minimum. In addition, while initial resource allocations may be dictated by worst case requirements of the hose profiles associated with the endpoints, real time measurement of actual VPN traffic allows adaptive resizing of network resource allocations so that initially allocated network resources may be applied to other uses that would otherwise be wasted. Thus, network resource allocations for the VPN achieves significant gains over private line networks.

The VPN service provider may achieve further allocation efficiency by exploiting multiplexing possibilities across multiple VPNs and resource sharing possibilities with other non-VPN communication traffic. In this way, VPN customers not only are relieved

of the need for detailed network expertise, but also gain significant cost savings due to network resource utilization efficiencies not possible with private link networks.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention is described in detail with reference to the following figures
5 wherein like numbers designate like elements, and wherein:

Fig. 1 is an exemplary block diagram of a VPN emulating a private line network;

Fig. 2 is an exemplary block diagram of a hose VPN;

Fig. 3 is an exemplary diagram of a customer managed hose VPN;

Fig. 4 is an exemplary diagram of a VPN service provider managed hose VPN;

10 Fig. 5 is an example of a set of network nodes connecting endpoints of a VPN;

Fig. 6 is a flow chart of an exemplary process for allocating and initializing a hose
VPN;

Fig. 7 is an exemplary block diagram of an access point;

15 Fig. 8 is a flow chart of an exemplary access point process for handling data
packets from a hose; and

Fig. 9 is a flow chart of an exemplary process for resizing allocated network
resources of a hose VPN.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Fig. 1 shows an example of a VPN 150 operating in a network system 100 that
20 provides a private line network service in a network 150. Customer networks 110-140
are VPN endpoints that communicate with one another via the network 150. In the
network system 100, each customer network 110-140 must allocate a set of customer
communication links 151-156 between pairs of source-destination endpoints and specify
communication link parameters such as bandwidth, packet delay, etc.

25 In the above VPN, the customer networks 110-140 bear the burden for network
resource allocation and, thus, operators of the customer networks 110-140 must have
precise knowledge of the amount of traffic flowing among the endpoints to the VPN to
set desired bandwidth and quality of service (QOS). Once allocated, network resources
such as bandwidth of nodes within the network 150 that are assigned for each
30 communication link 151-156 cannot be used for other communications even when not in
use for the intended VPN. Thus, inefficiencies in network resource utilization results

which reduces an ability of the network 150 to support a volume of communication that might otherwise be supported.

Fig. 2 is an exemplary diagram of a hose VPN system 200 that includes hoses 210-216 to a VPN and a network 250. The network may be of any kind such as a 5 switched network, an asynchronous transfer network (ATM), or other data or analog networks. For the following discussion, the Internet Protocol (IP) network is used as an example.

The IP network 250 includes a plurality of access points 218-224 where each of the access points 218-224 is associated with one of the hoses 210-216. The hose VPN 10 system 200 provides VPN services to customer networks 202-208 and each of the customer networks 202-208 gains access to the VPN via one of the hoses 210-216, for example.

The customer networks 202-208 may be any customer proprietary network such as a local area network (LAN) (e.g., an Internet Service Provider (ISP)), wide area 15 network (WAN), or the like. While Fig. 2 shows customer networks 202-208, any communication device capable of establishing a communication connection with the access points 218-224 may be used in place of or in conjunction with the customer networks 202-208 without departing from the spirit and scope of the invention.

The customer networks 202-208 access the VPN (i.e., utilize VPN service) 20 through a respective access point 218-224. The access points 218-224 are coupled to each other via communication links within the IP network. The access points 218-224 may be any type of device or system that provides access to the IP network 250. For example, the access points 218-224 may be Internet Service Providers (ISPs), network servers, routers, or the like.

A hose 210-216 provides access to a VPN that is supported by the IP network 250. When associated with the same VPN, the hoses 210-216 are endpoints of the VPN. Thus, each of the customer networks 202-208 communicates via one hose 210-216 to any 25 of the other endpoints (i.e., other customer networks 202-208 that are also part of the same VPN). Each hose 210-216 is specified by a service level agreements (SLA) without reference to other endpoints. Once agreed, the VPN guarantees performance of the hose 30 210-216 independent of a type or destination of communication traffic of the hose 210-

216 as long as the communication traffic remain within the SLA. Thus, a customer network 202-208 (or an operator of the customer network 202-208) need not specify performance for each of the endpoints that may be a destination. Rather, the customer network 202-208 only provide the SLA for its hose 210-216 and the VPN performs the functions required to meet the SLA.

5 Accordingly, once SLAs for all hoses 210-216 of a VPN are completed, the VPN is specified and the IP network 250 guarantees the performance such as connectivity and bandwidth required to meet the SLAs. In this way, operators of the customer networks may be relieved of tasks associated with routing, network management, network 10 performance, etc. In addition, the IP network 250 may be optimized to utilize resource sharing by multiplexing communication paths within each VPN and across multiple VPNs as well as take advantage of available resources supporting non-VPN communications.

15 SLAs for each of the hoses 210-216 may be based on one or more QOS requirements. For example, the hose 210 may have an SLA that requires three QOS levels of: QOS A, QOS B and QOS C. In addition, a time schedule may be associated with each of the QOS levels so that the level of support may be varied with time. For example, QOS A may require 10 megabits per second (mb/s) constant bit rate to support real time video communication, QOS B may require support for 5 mb/s file transfer 20 protocol (ftp) data transfers while QOS C may require 56 kilobits per second (kb/s) for e-mail. Also, QOS A may be needed only during working hours from 8 a.m. to 5 p.m., Monday through Friday, while QOS B may be needed for afternoon and evening hours between 1 p.m. and 8 p.m. and QOS C may be always required. All of the above requirements are collected together as a hose profile corresponding to each hose 210-216.

25 The VPN service provider guarantees that the hose profile is met for each of the hoses 210-216 of the VPN.

30 The SLA for hoses 210-216 may be based on traffic characteristics such as reasonable delay, packet loss rates, jitter, bandwidth minimums, and the like for various time periods. For example, an IP voice VPN service provider might require tight bounds on the per-packet loss rates, delay and possibly the amount of jitter. On the other hand, a data only VPN service provider might have relatively less stringent delay requirements.

Also, as mentioned earlier, these requirements may be needed only for certain time intervals and may be relaxed during other time intervals.

The interface between the customer networks 202-208 and the associated hoses 210-216 may be managed by either the customer or the VPN service provider. If the 5 VPN is established for a single QOS, then the management task is reduced to guaranteeing an effective bandwidth among all the access points 218-224 of the VPN. While QOS may be specified in terms of bandwidth, end-to-end delay, bit rate, error rate, etc., all of these qualities may be reduced to an effective bandwidth which the VPN service provider guarantees for each of the hoses 210-216.

10 For example, if a VPN is established between customer networks 202 and 208 and the associated hoses 210 and 216 are specified to require a 10 mb/s for each direction, the VPN service provider may satisfy the above VPN requirements by allocating 10 mb/s bandwidth in the access points 218 and 224 as well as 10 mb/s duplex connections between the access points 218 and 224. Assuming that a shortest path between the access points 218 and 224 traverses three nodes of the IP network, the VPN service provider 15 must allocate 10 mb/s of bandwidth in each of these nodes for both directions.

20 While the above simple VPN may appear to be supported by forming a pipe between the access points 218 and 224 without any multiplexing opportunities, the VPN service provider may take advantage of resource sharing between the resources allocated for the VPN and resources devoted for other non-VPN communications. For example, while 10 mb/s may be initially allocated for each direction to support a VPN, the network resource allocations may be resized based on actual VPN traffic experienced between the access points 218 and 224. In this way, the IP network resources may be adaptively 25 allocated so that un-utilized resources may be devoted for other IP network communications.

30 The VPN management task becomes much more difficult when multiple QOS levels are desired. For these more complex management situations, the management of the customer network interface may be divided into four cases as shown in Table 1 below. Cases 1 and 2 correspond to customer managed hoses 210-216 and cases 3 and 4 correspond to VPN service provider managed hoses 210-216.

In case 1, the customer may manage the outgoing hose traffic by controlling the rate of data packets output through the hoses 210-216 from different applications based on a priority scheme. Thus, higher priority applications may output more data packets than lower priority applications. After the data packets enter the hoses 210-216, the data packets are not distinguished from each other but are only identified by a VPN identification (ID), for example. Thus, for case 1, the VPN service provider does not distinguish one data packet from another within a single VPN, but only distinguishes one VPN from another. Thus, for case 1, the VPN service provider may only take advantage of potential multiplexing opportunities based on a single effective bandwidth hose profile. Other techniques may also be used such as weighted fair queuing so that different applications within the customer network may be assigned different weights and data packets from each application are allocated bandwidths through the hoses 210-216 based on the associated weights.

15

Table 1
Customer Network Interface

	Customer Managed	VPN Service Provider Managed One VPN for each QOS
Hose Input Scheduling	<u>Case 1</u> Customer controls the volume of packets from applications based on priority or weighted fair queuing techniques, for example.	<u>Case 3</u> Customer assigns applications to appropriate VPNs. The VPN service provider schedules customer communication traffic based on the QOS and VPN.
Packet Marking	<u>Case 2</u> Customer marks data packets. Network processes marked data packets up to the hose profile.	<u>Case 4</u> Customer and VPN service provider mark packets. Customer markings are honored as long as hose profile is met. Can optimize across VPNs having same QOS within the network.

20

In case 2, the customer manages the communication traffic through the hoses 210-216 by marking individual data packets. For example, each data packet is marked corresponding to the QOS assigned to the corresponding application. All the marked data packets are output to the hoses 210-216 and the VPN service provider processes the data packets based on the mark of each data packet. The VPN service provider will process all the marked data packets up to the hose profile. Data packets output through the hoses

210-216 that exceed the hose profile, may be dropped or handled in a best effort manner which may be specified in the hose profile.

Fig. 3 shows a diagram illustrating a VPN having customer managed hoses. Each of the customer networks 202-208 may independently define the hose profile for its respective hose 210-216. For example, the customer network 202 has a hose profile that has requirements for QOS A, QOS B and QOS C data packets while the hose profile for the hose 216 specifies requirements for QOS A and QOS B data packets. Similarly, a hose profile for the hose 212 specifies requirements for QOS A and QOS C data packets while a hose profile for the hose 214 specifies requirements for QOS B and QOS C data packets. If the customers use hose input scheduling, the desired communication performance is managed by the customer networks 202-208 allocating the hose bandwidth to its applications. The VPN guarantees the aggregate effective bandwidth specified in each of the hose profiles. If data packet marking is used as in case 2, the data packets are marked as QOS A, QOS B or QOS C and the VPN service provider handles the data packets according to the effective bandwidth specified in the hose profiles.

In case 3, the VPN service provider manages the hoses by defining a separate VPN for each QOS level. The customer networks 202-208 merely assign the appropriate VPN to applications that require the QOS level guaranteed by that VPN. Thus, the customer networks 202-208 do not perform any scheduling but rather output data packets of different QOS levels to respective VPNs.

In case 4, the VPN service provider establishes various VPNs corresponding to the QOS levels required and the customer networks 202-208 mark each of the data packets corresponding to the QOS level. In this case, the VPN service provider honors the markings of the data packets until the hose profile is reached. For example, if the customer network 202 marks a data packet to be QOS A, the VPN service provider transmits the data packet in the QOS A VPN until the QOS A hose profile is reached. QOS A marked data packets that exceed the hose profile may be given lower priority or best effort performance depending on network conditions. In some situations, the hose profile may specify burst traffic conditions where for a prespecified amount of time, the bandwidth may exceed a base level by some predetermined amount. In this situation, QOS level markings of data packets are ignored only when the burst condition parameters

of the hose profile are exceeded. Thus, the VPN service provider is policing the flow to ensure that marking apply only to communication traffic within the hose profile.

Fig. 4 shows a VPN that corresponds to the VPN shown in Fig. 3 with the exception that the VPN shown in Fig. 4 is a service provider managed VPN. Because the customer network 202 requires three QOS levels, the hose 210 includes three sub-hoses, one sub-hose corresponding to each of the QOS levels A, B and C. Similarly, the other hoses 212-216 may include sub-hoses for the corresponding QOS levels specified by each of the respective customer networks 204-208.

The VPN shown in Fig. 3 provides benefits for the customer networks 202-208 because the specification of the hose profile is simplified. The simplification is achieved because the customer networks 202-208 are required only to specify aggregate hose profiles instead of a specific hose profile for each QOS level. However, because only the aggregate hose profile is provided, the VPN service provider has less opportunity to take advantage of possible multiplexing than if more specific hose profile specifications were provided. Thus, the VPN service provider may be forced to allocate (at least initially) the maximum required bandwidths to guarantee the hose profile required performances.

In contrast, in the VPN shown in Fig. 4, the customer provides the VPN service provider bandwidth information required for each of the desired QOS levels for each of the endpoints. Thus, multiplexing among endpoints of each VPN may be achieved as well as resource sharing may be exploited across the VPNs allocated to all the QOS levels to reduce IP network resource allocations.

The VPN service provider may allocate network resources based on network optimization options as shown in Table 2 below. There are three possible network optimization strategies: 1) pipe; 2) source/sink tree; and 3) total VPN. As shown in Table 2, the degree of multiplexing ranges from low to high as discussed below.

Table 2
Network Optimization Options

Network Strategy	Degree of Multiplexing
Pipe	Low
Source/Sink Tree	Medium
Total VPN	High

Fig. 5 shows an example of a virtual network having customer networks 202, 206 and 208 connected to hoses 210, 214 and 216 as endpoints. For discussion purposes, assume that nodes 302 labeled A-G are allocated for the above VPN.

Table 3 below shows the pipe network optimization in terms of the nodes A-G.

5 Assume that the hose profile corresponding to each of the hoses 210-216 requires 1 mb/s for communication between the customer network 202 to the customer network 208 and between the customer network 202 and the customer network 206. If the shortest routing path between the customer networks 202 and 208 is via nodes A-C-B-E then the total bandwidth allocation in the network may be represented by the sum of the bandwidths between each adjacent node. Thus, the bandwidth between nodes A and C is 1 mb/s; the bandwidth between nodes C and B is 1 mb/s; and the bandwidth between nodes B and E is 1 mb/s. Thus, the network resource allocation between the customer network 202 and customer network 208 is 3 mb/s.

10
15 Table 3
Pipe Network Optimization

Endpoints	Bandwidth	Shortest Routing Path	Allocation	Total Allocation
A → E	1 mb/s	A→C→B→E	3 mb/s	6 mb/s
A → G	1 mb/s	A→C→D→G	3 mb/s	

Similar to the above, the network resource allocation between the customer networks 202 and 206 is 3 mb/s for the routing path of A-C-D-G. Thus, the total network resource allocation to support the hose profile for the hose 210 is 6 mb/s.

20 The above pipe network optimization basically allocates the maximum bandwidth required as if a pipe is constructed between nodes A and E and between nodes A and G. While initial bandwidth allocation may be 6 mb/s, the VPN service provider may monitor the actual VPN traffic and "resize" the network resource allocations based on the monitoring result. In this way, the network resource allocation may be adjusted to track 25 the actual bandwidth usage for the VPN and adaptively adjust the network allocations accordingly.

In addition, the VPN service provider may share network resources with other non-VPN communication traffic so that when network resources allocated to the VPN are not used, the VPN service provider may utilize the allocated network resources for other

purposes. For example, the unused network resources may be used for best effort IP network traffic until those resources are required to guarantee VPN performance.

Table 4 below shows an example of a first source tree network optimization. A source tree is a tree of connecting paths that emanate from a source hose 210-216 to a destination hose 210-216. For example, the source tree for the hose 210 originates from node A and extends to node C. If the shortest routing path is assumed, the source tree branches at node C to nodes B and D. The node B branch continues to node E while the node E branch continues to node G. This source tree obviates the fact that the branch A-C is common between the two branches C-B and C-D. Because the common branch A-C is only required to support the bandwidth allocated to the source, the hose 210, the bandwidth allocation for the branch A-C does not need to exceed the hose profile of the hose 210. Thus, as shown in Table 4, the network resource allocation may be 3 mb/s for the source tree branches A-C-B-E and 2 mb/s for the branches C-D-G. The total bandwidth allocation is 5 mb/s which is 1 mb/s less than the pipe network optimization of Table 3. Thus, the bandwidth gain of the source tree network optimization over the pipe network optimization is 6/5 or 1.2.

Table 4
First Source Tree Network Optimization

Endpoints	Bandwidth	Shortest Routing Path	Allocation	Total Allocation
A → E	1 mb/s	A → C → B → E	3 mb/s	5 mb/s
A → G	1 mb/s	→ D → G	2 mb/s	

Table 5 below shows a second source tree network optimization that takes advantage of an explicit routing path instead of relying on the shortest routing path. For example, instead of routing data packets through branches C-B-E and C-D-G, the data packets are routed through branches C-F-E and C-F-G, the branch between nodes C and F become common for the tree source between the hose 210 to the hose 216 and the hose 210 to the hose 214. Since the branches A-C and C-F of the source tree are required to support only the hose profile for hose 210, the bandwidth allocation for the branches A-C and C-F are 1 mb/s each. The remaining allocations are for the branches F-E and F-G which are both limited to the hose profiles for the hoses 216 and 214, respectively, and are both 1 mb/s. Thus, the total bandwidth allocation for the second source tree network

optimization is 4 mb/s. Accordingly, the gain of the second source tree network optimization over the pipe network optimization is 6/4 or 1.5.

Table 5
Second Source Tree Network Optimization

Endpoints	Bandwidth	Explicit Routing Path	Allocation	Total Allocation
A → E	1 mb/s	A → C → F → E	3 mb/s	4 mb/s
A → G	1 mb/s	→ G	1 mb/s	

5

The above source tree network optimizations take advantage of the ability to multiplex branches of the source tree for communication traffic destined to different endpoints. Thus, the branch A-C is multiplexed between the communication traffic between the hose 210 to the hose 216 and between the hose 210 and the hose 214.

10 Similarly, the second source tree network optimization multiplexes both the A-C branch and the C-F branch between the hose 210 to the hose 216 and the hose 210 to the hose 214 communication traffic.

15 Table 6 below shows a total VPN allocation for the second source tree optimization where all possible VPN communication traffic is included. The first row of table 6 shows the communication traffic between the hose 210 and the hoses 214 and 216, the second row shows the communication between the hose 216 and the hoses 210 and 214 and the third row shows the communication between the hose 214 and the hoses 210 and 216. As before, if the explicit routing path of the above second source tree network optimization is used, the total network allocation is 4 mb/s for each of the above hoses 20 210-216 which results in a total network allocation of 12 mb/s.

Table 6
Total VPN Allocation for Second Source Tree Optimization

Endpoints	Bandwidth	Explicit Routing Path	Allocation	Total Allocation
A → E	1 mb/s	A → C → F → E	3 mb/s	12 mb/s
A → G	1 mb/s	→ G	1 mb/s	
E → A	1 mb/s	E → F → C → A	3 mb/s	
E → G	1 mb/s	→ G	1 mb/s	
G → A	1 mb/s	G → F → C → A	3 mb/s	
G → E	1 mb/s	→ E	1 mb/s	

25

Table 7 below shows a VPN network optimization where further multiplexing may be exploited. If the network resource allocation is first performed for satisfying the hose profile for the hose 210 (row 1), then path A-C-F-E is allocated for communication

traffic between the hose 210 and the hose 216 and an additional allocation for path F-G is needed for communication traffic between the hose 210 and the hose 214. This allocation is identical with the allocation for the second source tree network optimization which takes advantage of multiplexing along paths A-C and C-F. In row 2, the allocation for the communication traffic from the hose 216 to the hose 210 requires 1 mb/s for each of the paths E-F-C-A. The allocation for the communication traffic between the hose 216 and the hose 214 requires the paths E-F and F-G. However, the path E-F may be multiplexed between the communication traffic of the hose 216 to the hose 210 and the hose 216 to the hose 214. Similarly, the path F-G may be multiplexed between the traffic from the hose 216 to the hose 214 and the hose 210 to the hose 214 as shown in row 1. Thus, the allocation for the path between F to G that was allocated in the second source tree optimization is not necessary. Those not allocated are shown in bold and italicized.

Table 7
VPN Network Optimization

Endpoints	Bandwidth	Explicit Routing Path	Allocation	Total Allocation
A → E	1 mb/s	A → C → F → E	3 mb/s	8 mb/s
A → G	1 mb/s	→ G	1 mb/s	
E → A	1 mb/s	E → F → C → A	3 mb/s	
E → G	1 mb/s	→ G		
G → A	1 mb/s	G → F → C → A	1 mb/s	
G → E	1 mb/s	→ E		

As shown in row 3, further multiplexing may be achieved for the communication traffic between the hose 214 and the hoses 210 and 216. Aside from the path G-F that is required to guarantee the hose profile corresponding to the hose 214, the remaining paths may be multiplexed between the communication traffic from the hose 214 to the hose 210 and the communication traffic from the hose 216 and the hose 210 as well as the communication between the hose 210 and the hose 216. Thus, the additional bandwidth allocation to support the hose profile corresponding to the hose 214 is only 1 mb/s. Accordingly, the total resource allocation for the VPN network optimization is 8 mb/s. Thus, the VPN network optimization has a gain of 3 over the pipe network optimization.

Fig. 6 shows a flow chart for an exemplary process to initially allocate network resources for a hose VPN. In step 1000, a request is received for VPN service and the

process goes to step 1002. In step 1002, it is determined whether allocation of hose resources should be customer managed or VPN service provider managed. If customer managed, the process goes to step 1004; otherwise, the process goes to step 1008. In step 1004, the VPN service provider negotiates hose profiles with the customer (or with each of the customer networks 202-208). Such negotiations may be conducted automatically via currently available or future protocols or may be negotiated between the VPN service provider business entity and the customer business entity. The negotiations may consider tradeoffs between QOS desired and costs incurred, for example. After the hose profiles are determined, the process goes to step 1005. In step 1005, the VPN service provider assigns a VPN identification (ID) to the requested VPN and goes to step 1006. In step 1006, the VPN service provider determines optimal routing to guarantee the required services as specified in the hose profiles and goes to step 1012.

In step 1008, the VPN service provider negotiates hose profiles for each of the QOS classes desired and goes to step 1009. In step 1009, the VPN service provider assigns VPN IDs for each of the QOS classes and goes to step 1010. In step 1010, the VPN service provider determines the optimal routing for each of the VPNs. For example, the VPN service provider may consider any multiplexing that may be achieved for each VPN and/or across all the VPNs. After the optimal routing has been determined, the process goes to step 1012.

In step 1012, the process allocates the network resources based upon the optimal routing determined in prior steps and goes to step 1014. In step 1014, the process determines whether the customer has chosen to mark the data packets that are transmitted. If packet marking is desired, the process goes to step 1016; otherwise, the process goes to step 1018. In step 1016, the VPN service provider receives the packet marks chosen by the customer (or the VPN service provider provides the packet marks to the customer) and initializes the packet marks in the network nodes that are selected based on the optimal routing paths and goes to step 1018. In step 1018, the VPN service provider initializes the allocated network resources and commences the VPN process and goes to step 1020 and ends the initial network resource allocation process.

The VPN service provider may monitor VPN traffic to predict traffic rates so that SLAs for hoses 210-216 may be adjusted or resized to actual usage patterns. The access

points 218-224 and/or other network nodes may monitor the incoming traffic to hoses 210-216 to collect hose traffic information as well as to ensure that the hose profiles are not violated. Similarly, traffic at a hose egress (i.e., traffic that has traversed the network) may be monitored for the same purposes.

5 Based on the information that is collected by the above monitoring, the VPN service provider may use network resource reservation mechanisms to resize current network resource allocations and provide prediction data for resizing the hose profile. Initially, the network resource allocations may be made statically by taking into account worst case resource demands. The initial allocations may then be resized dynamically 10 based on monitoring measurements. The resizing of the allocation may be restrained by the hose profiles as specified in the SLAs such that the resizing cannot cause the hose capacity to fall below a level required to provide the required QOS requirements. However, when within the hose profiles, the resizing of the resource allocation allows 15 more resources to be freed-up for use by other traffic and thus, more efficient use of the resources is possible. In addition, the hose profiles themselves may also be renegotiated based on the usage data generated based on the data produced by the monitoring process. The monitoring and resizing functions are discussed below in connection with an exemplary access point block diagram.

Fig. 7 is an exemplary block diagram of an access point 218, for example. 20 Although the access point 218 is described, it should be appreciated that the other access points 220-224 may have similar architecture. In addition, with the exception of the hose interface 330, Fig. 7 may also represent a node in the IP network 250. The access point 218 includes a controller 310, an IP network interface 320, a customer network interface 330, a resource usage measurement device 340, a memory 350 and a prediction device 360. These elements are coupled together via a signal bus 370. While a bus architecture 25 is shown for ease of illustration, other architectures are possible, as is well known to one of ordinary skill.

When the customer network 202 requests a hose 210 for accessing a VPN, the 30 access point 218 negotiates with the customer network 202 to determine an SLA to specify a hose profile for the hose 210. The successfully negotiated SLA is stored in the memory 350. The hose profile may include information related to the amount of traffic

expected and the QOS(s) desired. The QOS may be a measure of data packet loss, delay, jitter and the like as discussed earlier. Other SLA parameters may be stored in the memory 350 and used to determine resizing of hose capacity as described hereafter.

Based on the hose profile, the access point 218 establishes the hose 210 between the customer network 210 and the VPN. The access point 218 may also communicate with other nodes of the IP network 250 to establish routing paths. As discussed earlier, various VPN network optimization schemes may be used. Thus, the access point 218 may perform management functions of the VPN service provider to establish and administer the VPN.

The traffic information transmitted over the hoses 210-216 may be data packets. Each data packet may include header information identifying the source, destination, data packet sequence number, QOS markings and the like. Based on the header information of the data packet, the controller 310 forwards the data packet via the IP network interface 320 to a next node or access points 220-224 in the IP network which may be determined by a routing scheme selected by the VPN service provider, such as shortest path or explicit routing path.

The access point 218 may perform monitoring functions by measuring current resource usage via the resource usage measurement device 340. For example, the resource usage measurement device 340 may measure the used bandwidth, packet loss, and/or other transmission parameters via direct measurement of the data stream, a standard network message protocol (SNMP) or other signaling interfaces such as the SS7. The resource usage measurements are then input to the prediction device 360 and used to predict hose capacity needed for the hose 210. Based on the predicted hose capacity, the controller 310 may renegotiate the capacity of the hose 210 to resize the hose 210 for greater network resource allocation efficiency. Other network resource allocations may also be dynamically resized based on the predictions generated by the prediction device 360.

As an example of a hose capacity prediction technique, traffic flow measurements gathered at regularly spaced instants during a window of duration T_{meas} may be considered. The measurements may be used to predict an effective bandwidth for the traffic flow over some window of duration T_{ren} following the measurement window.

The predicted effective bandwidth may be set to the maximum bandwidth measured during the measurement window T_{meas} . Alternatively, the predicted effective bandwidth may be set equal to $m + \alpha \sqrt{v}$, where m and v are respectively the mean and variance of the bandwidths sampled during the measurement window and α is a multiplier that controls the extent to which the negotiated bandwidth accommodates variability in the samples. The interpretation of α is that in a Gaussian approximation to the bandwidth distribution. The bandwidth $m + \alpha \sqrt{v}$ is expected to be exceeded with probability $1 - G(\alpha)$, where G is the cumulative distribution of the standard normal distribution.

If the interval between hose capacity renegotiations encompasses periods of systematic variation, the above predictions may become inaccurate. In order to remedy these inaccuracies, two approaches may be followed. First, a worst case predictor over the largest time scale of variation, e.g., the maximum bandwidth over a day for telephone traffic, may be used to compensate for inaccuracies. Second, historical data may be used to predict trends. For example, when average telephone call arrival rates are a known constant process $Q(t)$, then instead of using a predicted bandwidth $S(t)$ directly, the predictor $S(t)Q(t+T_{\text{ren}})/Q(t)$ may be used. Here, the ratio $Q(t+T_{\text{ren}})/Q(t)$ is used to model the systematic change of the arrival rate upwards or downwards.

In addition to the inaccuracies created by systematic variations, there are two statistical effects which, if uncorrected, may cause the prediction to underestimate bandwidth requirements. The first is sampling error and the second is short-time scale burstiness. Estimation of mean and variance is subject to inherent sampling error since the estimates are themselves random variable. This additional variability can lead to violation of target quality metrics if estimated parameters are assumed to be the true ones. The sampling error for n samples may be avoided by increasing the multiplier α to $\alpha' = ((1+n)(\varepsilon^{(\alpha^2)/n} - 1))^{1/2} > \alpha$. For example, if $n = 60$ and $\alpha = 3$, $\alpha' = 3.14$.

Burstiness at multiple time scales has been observed in Internet data traffic. The variability of the window averaged bandwidth of such traffic over a given window may increase for smaller window sizes. A predictor based on a given sampling window can underestimate the bandwidth required to satisfy the QOS of the SLA specified at shorter

time scales. A prior knowledge of the scaling relations between bandwidth variance at different time scales can be used to correct the multiplier α in order to accommodate short time scale variability.

Fig. 8 shows a flow chart for an exemplary process of an access point 218-224 such as access point 218 for supporting a VPN. In step 2000, the hose interface 330 receives data packets from the hose 210 and goes to step 2002. In step 2002, the controller 310 determines whether the hose profile has been exceeded. If exceeded, the process goes to step 2004; otherwise, the process goes to step 2006. In step 2004, the controller 310 forwards the received data packet to the IP network interface 320 to be transmitted to the destination via a best effort service, for example. The best effort service may include dropping the data packet because network resources are not available for sending the data packet to its destination. Then, the process goes to step 2008. In step 2006, the controller 310 examines the data packet header to determine the QOS of the data packet and forwards the data packet to the next node 302 based on a routing path that is initialized when the VPN was established. Then the process goes to step 2008. In step 2008, the controller 310 determines whether more data packets have been received via the hose interface 330. If more data packets have been received, the process returns to step 2000; otherwise, the process goes to step 2010. In step 2010, the controller 310 determines whether a system end condition is detected. If detected, the process goes to step 2012 and ends; otherwise, the process returns to step 2008.

Fig. 9 shows a flow chart for a monitoring process of the VPN service provider. In step 3000, the VPN service provider determines whether the time to determine whether resizing is necessary has been reached. If reached, the VPN service provider goes to step 3002; otherwise, the VPN service provider returns to step 3000. In step 3002, the VPN service provider collects measurement data from nodes 302 and/or access points 218-224 and optionally executes the traffic predictor via the prediction device 350, for example. Then the VPN service provider goes to step 3004. In step 3004, the VPN service provider compares the predicted traffic to allocated thresholds.

Multiple thresholds may be used. For example, when the predicted capacity exceeds an upper bound threshold, then additional capacity should be allocated for the VPN and the VPN service provider goes to step 3008. When the predicted capacity falls

below a lower bound threshold, then the allocated resources should be reduced and the VPN service provider goes to step 3006. If the predicted capacity is between the upper bound and the lower bound thresholds, then the current resource allocation is acceptable and the VPN service provider goes to step 3022.

5 In step 3006, the VPN service provider reduces resource allocation of access points 218-224 and nodes 302 that are within the routing paths of the VPN and the VPN service provider goes to step 3011. In step 3011, the VPN service provider determines whether the hose profiles of the VPN should be renegotiated. If renegotiation is required, the VPN service provider goes to step 3013; otherwise, the VPN service provider goes to 10 step 3022.

In step 3008, the VPN service provider determines whether the predicted capacity exceeds the hose profile thresholds. For example, thresholds may be set so that when the predicted capacity is within a predetermined value specified in the hose profile, then an alarm condition is set to determine whether the hose profile is required to be renegotiated. 15 If the hose profile thresholds have been exceeded, the VPN service provider goes to step 3012; otherwise the VPN service provider goes to step 3010. In step 3012, the VPN service provider renegotiates the hose profiles to increase bandwidth and goes to step 3022.

20 In step 3010, the VPN service provider determines whether there is additional bandwidth within the nodes 302 and/or access points 218-224 of the routing path. If additional bandwidth is available, the VPN service provider goes to step 3016; otherwise, the VPN service provider goes to step 3014. In step 3016, the VPN service provider increases resource allocations based on the predictor capacities and goes to step 3022.

25 In step 3014, the VPN service provider determines whether additional bandwidth may be available in other nodes that may be a part of the routing path for the VPN. If available, the VPN service provider goes to step 3020; otherwise, the VPN service provider goes to step 3018. In step 3020, the VPN service provider changes the routing path of the VPN and allocates new resources so that the hose profile requirements may be met and goes to step 3022. In step 3018, the VPN service provider sets flags to indicate 30 that network resource upgrades may be required due to excess demand for VPN services. Then the VPN service provider goes to step 3022. In step 3022, the VPN service

provider determines whether a system end condition has been detected. If detected, the VPN service provider goes to step 3024 and ends the process; otherwise, the VPN service provider returns to step 3000.

The above described VPN service provider functions for establishing, managing and executing a VPN may be implemented by using a variety of available protocols such as the Integrated Service (IntServ), differentiated services (DiffServ), or Multi-Protocol Label Switching (MPLS) as defined in the Internet Engineering Task Force. For example, in an IntServ framework, a signaling protocol may be used to allocate resources for hoses 210-216 between selected endpoints on an as-needed basis. For this example, a customer network 202-208 may send signals via a signaling protocol to an access point 218-224 to establish a hose 210-216 and to identify endpoints of a VPN. The VPN service provider may respond by reserving resources within the IP network 250 by determining routing paths to the selected endpoints and setting up hoses 210-216 at each of the endpoints. The RSVP protocol may be used to maintain the resource reservations across the nodes of the established routing path. If the VPN that is established is no longer needed, the signaling via the RSVP protocol may be terminated resulting in releasing of the resources along the established routing path.

The DiffServ framework supports data packet marking so that various QOS levels may be implemented. The MPLS environment uses labels to place in the headers of data packets to determine routing paths that may be characterized by a sink tree. A sink tree is routed at a hose egress point and the branches extend to ingress points of other hoses 210-216. Such a sink tree may also be considered a label switch path (LSP) tree. For example, the VPN service provider may use traffic measurements to determine an actual spread of traffic from a hose egress point to several other hose ingress points and signal on each LSP involved to reserve the required resources. The signaling to create LSPs and to reserve reservations on such paths might be combined in a single protocol.

While this invention has been described with specific embodiments thereof, it is evident that many alternatives, modifications, and variations will be apparent to those skilled in the art. Accordingly, the preferred embodiments of the invention as set forth herein are intended to be illustrative, not limiting. Various changes may be made without departing from the spirit and scope of the invention.

WHAT IS CLAIMED IS:

1 1. A method for providing a virtual private network service, comprising:
2 establishing a hose for each of a plurality of endpoints of a virtual private
3 network;

4 coupling the hose to endpoints associated with other hoses via routing
5 paths in a network; and

6 allocating network resources to support communications between the hose
7 and the other hoses.

1 2. The method of claim 1, wherein the establishing comprises specifying a
2 service level agreement for the hose, the service level agreement including a hose profile
3 and other information for controlling and managing the hose.

1 3. The method of claim 2, wherein the establishing comprises:
2 first selecting one of a user managed hose type or a virtual private network
3 service provider managed hose type for the hose; and

4 second selecting whether to transmit marked data packets to the hose,
5 results of the first and second selecting steps being stored in the hose profile.

1 4. The method of claim 3, wherein if the user managed hose type is selected,
2 the method further comprising:

3 specifying one or more aggregate bandwidths for the hose; and
4 specifying a time schedule for each of the aggregate bandwidths, the
5 aggregate bandwidths and the time schedule being stored in the hose profile.

1 5. The method of claim 4, wherein if data packet marking is selected, the
2 method further comprising:

3 receiving information regarding data packet markings and a quality of
4 service corresponding to each of the data packet markings; and

5 initializing the allocated network resources to provide the quality of
6 service based on the data packet markings if conditions in the hose profile is not violated.

1 6. The method of claim 3, wherein if the virtual service provider managed
2 hose type is selected, the method further comprising:

3 receiving one or more quality of service levels for the hose;

establishing one or more sub-virtual private networks, each sub-virtual network corresponding to one of the quality of service levels;

specifying one or more bandwidths for the hose corresponding to each of the sub-virtual private networks; and

specifying one or more time schedules for the bandwidths, the bandwidths and the time schedules being stored in the hose profile.

7. The method of claim 6, wherein if data packet marking is selected, the method further comprising:

receiving information regarding data packet markings and a quality of service corresponding to each of the data packet markings; and

initializing the allocated network resources to provide the quality of service based on the data packet markings if conditions in the hose profile is not violated.

8. The method of claim 2, further comprising:

measuring communication traffic of allocated network resources to generate monitoring data;

generating a resizing condition based on the monitoring data; and

resizing the allocated network resources if the resizing condition is within one or more thresholds of the hose profile.

9. The method of claim 8, wherein the monitored data includes historical data, the method further comprising generating trend data to predict virtual private network usage.

10. The method of claim 8, wherein the resizing condition is one of above an upper bound threshold, below a lower bound threshold, and within the upper bound and lower bound thresholds.

11. The method of claim 10, further comprising:

reducing the allocated network resources if the resizing condition is below the lower bound threshold; and

increasing the allocated network resources if the resizing condition is above the upper bound threshold.

1 12. The method of claim 10, wherein if the resizing condition is below the
2 lower bound threshold by a predetermined amount, the method further comprising

3 renegotiating the hose profile to change the service level agreement to be more consistent
4 with the monitored data.

1 13. The method of claim 10, wherein if the resizing condition is above limits
2 set by the hose threshold, the method further comprising renegotiating the hose profile to
3 change the service level agreement to be more consistent with the monitored data.

1 14. The method of claim 8, wherein the resizing condition determined based
2 on a prediction of future virtual private network usage.

1 15. The method of claim 1, wherein the routing paths is determined based on
2 one or more of:

3 network connectivity;
4 a hose identification; and
5 virtual private network identification.

1 16. The method of claim 15, further comprising selecting the routing paths
2 based on a shortest distance between pairs of endpoints of the virtual private network to
3 form a pe between the pairs of the endpoints.

1 17. The method of claim 15, further comprising:
2 selecting the routing paths based on a source tree or a sink tree for each of
3 the endpoints; and

4 minimizing a bandwidth allocation between nodes of the network by
5 maximizing sharing of same paths for branches of the source or the sink tree extending
6 between different ones of the endpoints.

1 18. The method of claim 15, further comprising:
2 selecting the routing paths based on source trees or sink trees
3 corresponding to all endpoints of the virtual private network; and
4 minimizing a bandwidth allocation between nodes of the network by
5 maximizing sharing of same paths for branches of the source or the sink trees extending
6 between different ones of the endpoints.

1 19. The method of claim 15, further comprising:
2 selecting the routing paths based on source trees or sink trees
3 corresponding to all endpoints of one or more virtual private networks; and

minimizing a bandwidth allocation between nodes of the network by maximizing sharing of same paths for branches of the source or the sink trees extending between different ones of the endpoints for all the virtual private networks.

20. The method of claim 1, wherein the network is an Internet Protocol Network.

21. A virtual private network in an network, comprising:

- a plurality of endpoints, each of the endpoints having a hose; and
- a plurality of routing paths in the network, the routing paths coupling the hose to endpoints associated with other hoses; and

a virtual private network service provider; the virtual private network service provider allocating network resources to support communications between the hose and the other hoses.

22. The virtual private network of claim 21, wherein the virtual private network service provider receives a service level agreement for the hose, the service level agreement including a hose profile and other information for controlling and managing the hose.

23. The virtual private network of claim 22, wherein the virtual private network service provider receives a first selection of one of a user managed hose type or a virtual private network service provider managed hose type for the hose and a second selection of whether to transmit marked data packets to the hose, results of the first and second selections being stored in the hose profile.

24. The virtual private network of claim 23, wherein if the user managed hose type is selected, the virtual private network service provider receives a specification for one or more aggregate bandwidths for the hose, and a specification for a time schedule for each of the aggregate bandwidths, the aggregate bandwidths and the time schedule being stored in the hose profile.

25. The virtual private network of claim 24, wherein if data packet marking is selected, the virtual private network service provider receives information regarding data packet markings and a quality of service corresponding to each of the data packet markings and initializes the allocated network resources to provide the quality of service based on the data packet markings if conditions in the hose profile is not violated.

1 26. The virtual private network of claim 23, wherein if the virtual service
2 provider managed hose type is selected, the virtual private network service provider
3 receives one or more quality of service levels for the hose, establishes one or more sub-
4 virtual private networks, each sub-virtual network corresponding to one of the quality of
5 service levels, receives a specification of one or more bandwidths for the hose
6 corresponding to each of the sub-virtual private networks, and receives a specification of
7 one or more time schedules for the bandwidths, the bandwidths and the time schedules
8 being stored in the hose profile.

1 27. The virtual private network of claim 26, wherein if data packet marking is
2 selected, the virtual private network service provider receives information regarding data
3 packet markings and a quality of service corresponding to each of the data packet
4 markings, and initializes the allocated network resources to provide the quality of service
5 based on the data packet markings if conditions in the hose profile is not violated.

1 28. The virtual private network of claim 22, wherein the virtual private
2 network service provider measures communication traffic of allocated network resources
3 to generate monitoring data, generates a resizing condition based on the monitoring data,
4 and resizes the allocated network resources if the resizing condition is within thresholds
5 of the hose profile.

1 29. The virtual private network of claim 28, wherein the monitored data
2 includes historical data, the method further comprising generating trend data to predict
3 virtual private network usage.

1 30. The virtual private network of claim 28, wherein the resizing condition is
2 one of above an upper bound threshold, below a lower bound threshold, and within the
3 upper bound and lower bound thresholds.

1 31. The virtual private network of claim 30, wherein the virtual private
2 network service provider reduces the allocated network resources if the resizing condition
3 is below the lower bound threshold, and increases the allocated network resources if the
4 resizing condition is above the upper bound threshold.

1 32. The virtual private network of claim 30, wherein if the resizing condition
2 is below the lower bound threshold by a predetermined amount, the method further

3 comprising renegotiating the hose profile to change the service level agreement to be
4 more consistent with the monitored data.

1 33. The virtual private network of claim 30, wherein if the resizing condition
2 is above limits set by the hose threshold, the method further comprising renegotiating the
3 hose profile to change the service level agreement to be more consistent with the
4 monitored data.

1 34. The virtual private network of claim 28, wherein the resizing condition
2 determined based on a prediction of future virtual private network usage.

1 35. The virtual private network of claim 21, wherein the routing paths is
2 determined based on one or more of:

3 network connectivity;
4 a hose identification; and
5 virtual private network identification.

1 36. The virtual private network of claim 34, wherein the virtual private
2 network service provider selects the routing paths based on a shortest distance between
3 pairs of endpoints of the virtual private network to form a pipe between the pairs of the
4 endpoints.

1 37. The virtual private network of claim 34, wherein the virtual private
2 network service provider selects the routing paths based on a source tree or a sink tree for
3 each of the endpoints, and minimizes a bandwidth allocation between nodes of the
4 network by maximizing sharing of same paths for branches of the source or the sink tree
5 extending between different ones of the endpoints.

1 38. The virtual private network of claim 34, wherein the virtual private
2 network service provider selects the routing paths based on source trees or sink trees
3 corresponding to all endpoints of the virtual private network, and minimizes a bandwidth
4 allocation between nodes of the network by maximizing sharing of same paths for
5 branches of the source or the sink trees extending between different ones of the endpoints.

1 39. The virtual private network of claim 34, wherein the virtual private
2 network service provider selects the routing paths based on source trees or sink trees
3 corresponding to all endpoints of one or more virtual private networks, and minimizes a
4 bandwidth allocation between nodes of the network by maximizing sharing of same paths

- 5 for branches of the source or the sink trees extending between different ones of the
- 6 endpoints for all the virtual private networks.

ABSTRACT OF THE DISCLOSURE

The invention provides apparatus and methods for a Virtual Private Network (VPN) in a network that offers a simple user interface for efficient utilization of network resources. The VPN is defined for a specified set of endpoints each of which is associated with a single "hose." A hose provides access to the VPN through an access point which may be a node of the network, for example. The hose is a single interface to the VPN for communication to all other endpoints of the VPN. The VPN achieves network resource allocation efficiency by exploiting resource sharing possibilities via multiplexing routing paths between endpoints and dynamic resource allocation techniques that permit real time resource allocation resizing. When a VPN is established with a VPN service provider, the routing paths between the endpoints of the VPN is optimized for multiplexing opportunities so that resource allocations between nodes along routing paths within the IP network is reduced to a minimum.




Fig. 1




Fig. 2




Fig. 3




Fig. 4




Fig. 5




Fig. 6




FIG. 7





Fig. 8

Fig. 9



IN THE UNITED STATES
PATENT AND TRADEMARK OFFICE

Declaration and Power of Attorney

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe I am an original, first and joint inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled VIRTUAL PRIVATE NETWORK, the specification of which is attached hereto.

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by an amendment, if any, specifically referred to in this oath or declaration.

I acknowledge the duty to disclose all information known to me which is material to patentability as defined in Title 37, Code of Federal Regulations, 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, 119 of any foreign application(s) for patent or inventors' certificate listed below and have also identified below any foreign application for patent or inventors' certificate having a filing date before that of the application on which priority is claimed:

None

I hereby claim the benefit under Title 35, United States Code, 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, 112, we acknowledge the duty to disclose all information known to us to be material to patentability as defined in Title 37, Code of Federal Regulations, 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

**Serial No. 60/113,497
Filed December 22, 1998
Serial No. 60/104,756
Filed October 19, 1998
Status: Provisional applications**

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

I hereby appoint the following attorney(s) with full power of substitution and revocation, to prosecute said application, to make alterations and amendments therein, to receive the patent, and to transact all business in the Patent and Trademark Office connected therewith:

Samuel H. Dworetzky (Reg. No. 27,873)
 Thomas A. Restaino (Reg. No. 33,444)
 Michele L. Conover (Reg. No. 34,962)
 Jose de la Rosa (Reg. No. 34,810)
 Benjamin S. Lee (Reg. No. 42,787)
 Robert B. Levy (Reg. No. 28,234)
 Alfred G. Steinmetz (Reg. No. 22,971)

I also appoint James A. Oliff (Reg. No. 27,075), William P. Berridge (Reg. No. 30,024), Kirk M. Hudson (Reg. No. 27,562), Thomas J. Pardini (Reg. No. 30,411), Edward P. Walker (Reg. No. 31,450), Robert A. Miller (Reg. No. 32,771), Mario A. Costantino (Reg. No. 33,565), and Caroline D. Dennison (Reg. No. 34,494) of OLIFF & BERRIDGE, PLC as associate attorneys, with full power to prosecute said application, to make alterations and amendments therein, and to transact all business in the U.S. Patent and Trademark Office connected therewith.

Please address all correspondence S. H. Dworetzky, AT&T Corp., P.O. Box 4110, Middletown, New Jersey 07748. Telephone calls should be made to (973) 360-8120.

Full name of 1st joint inventor: Nicholas G. Duffield

Inventor's signature Nicholas G. Duffield Date 10/18/99

Residence: Hoboken, Hudson County, New Jersey

Citizenship: Great Britain

Post Office Address: 933 Park Avenue, #5R
 Hoboken, New Jersey 07030

Full name of 2nd joint inventor: Albert G. Greenberg

Inventor's signature

Albert G. Greenberg

Date

Oct 18, 1999

Residence: Millburn, Essex County, New Jersey

Citizenship: United States

Post Office Address: 76 Chestnut Street
Millburn, New Jersey

Full name of 3rd joint inventor: Pawan Goyal

Inventor's signature Pawan Goyal Date 10/19/99

Residence: Mountain View, Santa Clara County, California

Citizenship: India

Post Office Address: 777 W. Middlefield Road, Apt. 83
Mountain View, California 94043

Full name of 4th joint inventor: Partho P. Mishra

Inventor's signature



Date

10-14-99


Residence: Summit, Union County, New Jersey

Citizenship: India

Post Office Address: 24 Beekman Terrace
Summit, New Jersey 07901

Full name of 5th joint inventor: Kadangode K. Ramakrishnan

Inventor's signature



Date Oct 18, 1999


Residence: Berkeley Heights, Union County, New Jersey

Citizenship: United States

Post Office Address: 9 Highland Circle
Berkeley Heights, New Jersey 07922

Full name of 6th joint inventor: Jacobus Erasmus van der Merwe

Inventor's signature



Date

Oct 18, 1999

Residence: Summit, Union County, New Jersey

Citizenship: South African

Post Office Address: 5 Middle Avenue
Summit, New Jersey 07901